# Predicting election trends with Twitter: Hillary Clinton versus Donald Trump

Alexandre Bovet, Flaviano Morone, Hernán A. Makse

*Levich Institute and Physics Department, City College of New York, New York, New York 10031, USA*

arXiv:1610.01587v2 [cs.SI] 7 Nov 2016

**Abstract**

Forecasting opinion trends from real-time social media is the long-standing goal of modern-day big-data analytics. Despite its importance, there has been no conclusive scientific evidence so far that social media activity can capture the opinion of the general population at large. Here we develop analytic tools combining statistical physics of complex networks, percolation theory, natural language processing and machine learning classification to infer the opinion of Twitter users regarding the candidates of the 2016 US Presidential Election. Using a large-scale dataset of 73 million tweets collected from June 1st to September 1st 2016, we investigate the temporal social networks formed by the interactions among millions of Twitter users. We infer the support of each user to the presidential candidates and show that the resulting Twitter trends follow the New York Times National Polling Average, which represents an aggregate of hundreds of independent traditional polls, with remarkable accuracy ($r = 0.9$). More importantly, the Twitter opinion trend forecasts the aggregated NYT polls by 6 to 15 days, showing that Twitter can be an early warning signal of global opinion trends at the national level. Our analytics [available at `kcorelab.com`] unleash the power of Twitter to predict social trends from elections, brands to political movements, and at a fraction of the cost of national polls. UPDATE November 7, 2016: We apply our analysis to the final part of the presidential election. Our Twitter opinion trend anticipates The Upshot NYT polling average by up to 13 crucial days and shows Clinton as the winner of the Twitter opinion by 55.5% versus Trump 44.5% of the votes (normalized to two candidates).

## 1 Introduction

Several works have showed the potential of social media, such as the microblogging platform Twitter, as a gauge for analyzing the public sentiment in general [1–4], stock markets or sales performance [6–8]. With the increasing importance of Twitter in political discussions, a number of studies [9–14] also investigated the possibility to predict political elections from Twitter data, sometimes with mixed results. Most approaches either compare the volume of tweets related to the different candidates with the election results or perform a sentiment analysis of the tweets to infer the global sentiment toward the candidates or political parties. Different methods exist to infer the sentiment polarity expressed in a tweet as positive or negative [15]. Most work studying the political sentiment in Twitter perform sentiment analysis with lexicon-based approaches [17–19], relying on collections of pre-compiled terms that are known to express a sentiment. However, Twitter language is known to be unstructured, informal and to contain spelling mistakes, urls, hashtags, emoticons and usernames that limits the applicability of predefined lexicons. A different approach for sentiment analysis is based on machine learning where a set of existing

labeled documents, from which features are extracted, is used to classify the rest of the documents [15]. For example, the presence of emoticons in the tweets was successfully used to label large training sets of positive and negative tweets [2, 20].
Here, we focus on the 2016 US Presidential Election by monitoring Twitter activity regarding the two top candidates to the presidency: Hillary Clinton (Democratic Party) and Donald J. Trump (Republican Party). We develop a supervised learning approach to classify the tweets as supporting or opposing the political candidates. Contrary to previous studies, we do not try to classify tweets as expressing positive or negative sentiment. Instead, we classify the tweets as supporting or opposing one of the candidates. We therefore avoid the problem of correctly assigning the object of the sentiment of a tweet. Indeed, a tweet containing a mention of Donald Trump and expressing a negative sentiment might be expressing opposition to Donald Trump as well as support. In this case, the context of the tweet is extremely important. Using an in-domain training set not only helps us to capture the informalities of Twitter language, it also permits us to capture the rich context of the 2016 US election.

## 2 Results

### 2.1 Building the network of Twitter users

We collect 73 million tweets using Twitter Search API from June 1st, 2016 to September 1st, 2016 mentioning the two top candidates to the 2016 US Presidential Election by using the following queries: *trump OR realdonaldtrump OR donaldtrump* and *hillary OR clinton OR hillaryclinton*. The total number of users in our dataset is 6.7 million with an average of about 290,000 distinct users per day.
We then build the daily social networks from user interactions following the methods developed in Ref. [21] (see methods 5.1). Using concepts borrowed from percolation theory [22, 23] we define different connected components to characterize the connectivity properties of the whole network: the strongly connected giant component (SCGC), weakly connected giant component (WCGC) and the corona (the rest of the network composed of smaller components that do not belong to the giant components SCGC and WCGC)
The SCGC is formed by the users that are part of interaction loops and are the most involved in discussions while WCGC is formed by users that not necessarily have reciprocal interactions with other users (see Fig. 1a). A typical daily network is shown in Fig. 1b. We monitor the evolution of the size of the SCGC, WCGC and the corona as shown in Fig. 2.
The size of the SCGC varies between approximately 15,000 and 35,000 users and is approximately 10 times smaller than the WCGC (see Fig. 2a). Fluctuations in the size of the three compartments are visible in the large spikes in activity occurring during important events that happened during the period of observation. For instance, on June 6, the Associated Press announced that Hillary Clinton had secured enough delegates to be the nominee of the Democratic Party. Bernie Sanders (who was the second contender for the Democratic Nomination) officially terminated his campaign and endorsed Hillary Clinton on July 12. The Republican and Democratic Conventions were held between June 18-21 and June 25-28, respectively. The spikes in activity related to these events are more

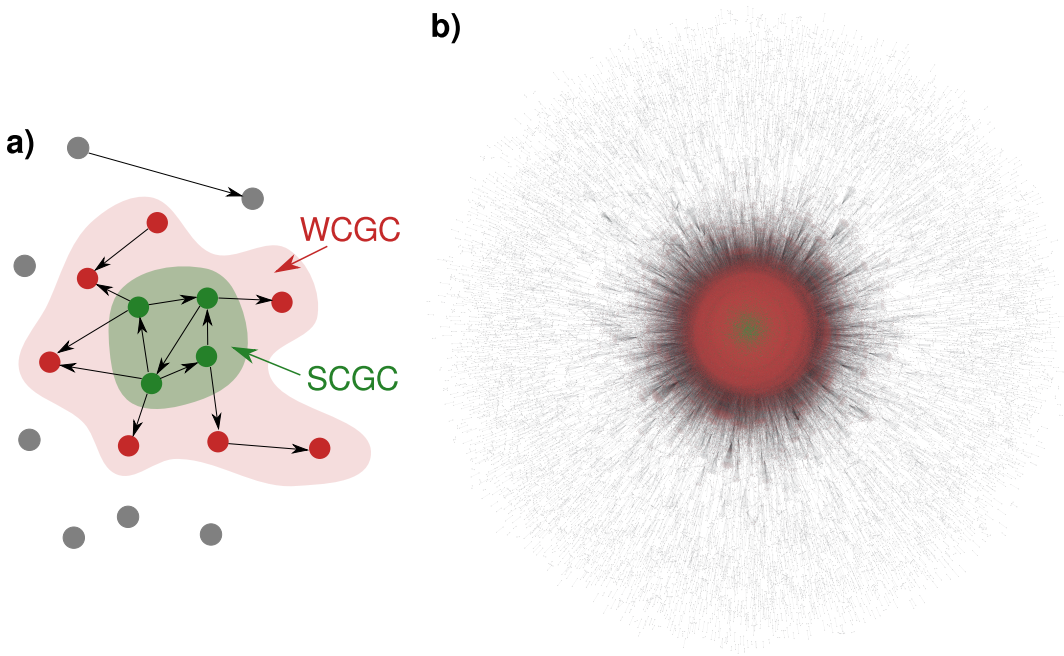


**Figure 1 | Definition of network components of users of the Twitter-election sphere.** (a) Sketch representing the weakly (red) and strongly (green) connected giant components and the corona (gray). (b) Visualization of a real influence daily network reconstructed from our Twitter dataset. The strongly connected giant component (green) is the largest maximal set of nodes where there exists a path in both directions between each pair of nodes. The weakly connected giant component (red) is the largest maximal set of nodes where there exists a path in at least one direction between each pair of nodes. The corona (grey) is formed by the smaller components.

important in the WCGC and Corona than in the SCGC. The number of new users, that arrive in our dataset for the first time, in each compartment is displayed in green in Fig. 2. Most of the new users arrive in the WCGC or the corona while relatively few users join directly the strongly connected component. This is expected as the users belonging to SCGC are those who are supposed to be the influencers in the campaigns, since for those people in the SCGC, the information can arrive from any other member of the giant component, and, viceversa, the information can flow from the member to any other user in the SCGC. Thus, it may take time for a new arrival to belong to the SCGC of influencers in the campaign. After the first week of observation, the number of new users arriving directly to the SCGC per day stays stable below 1,000.

## 2.2 Inferring the opinion of Twitter users

We use a set of hashtags expressing opinion to build a set of labeled tweets used to train a machine learning classifier (see methods 5.2). Figure 3 displays a force-layout of the network of hashtags discovered with our algorithm. Hashtags are colored according to the four categories, pro-Trump (red), anti-Hillary (orange), pro-Clinton (blue) and anti-Trump (purple). Two main clusters, formed by the pro-Trump and anti-Hillary on the top and pro-Clinton and anti-Trump on the bottom, are visible, indicating a strong relation between the usage of hashtags in these two pairs of categories. We identify more hashtags

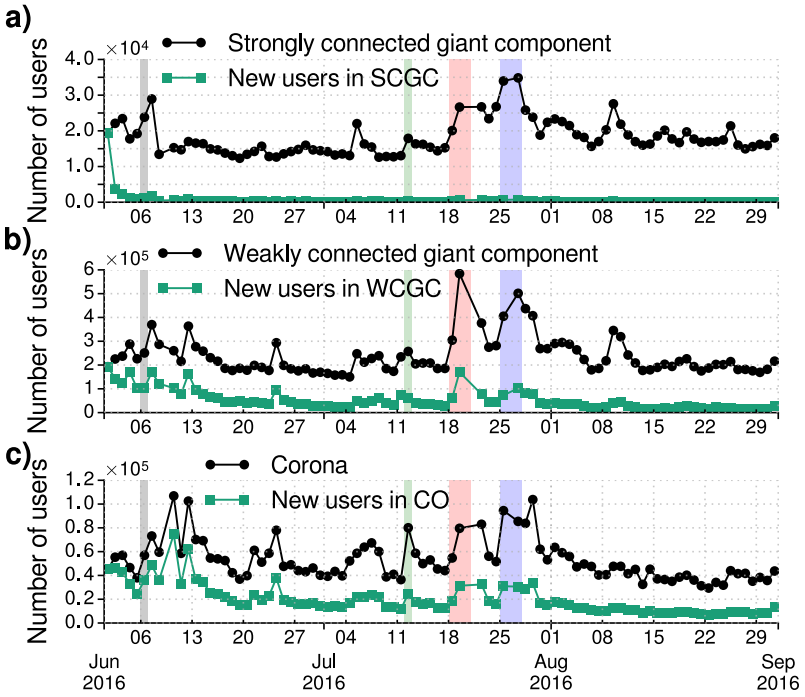


**Figure 2 | Temporal evolution of daily network components of Twitter election users.** (a) Total number of users in the daily strongly connected giant component versus time, (b) weakly connected giant component and (c) the corona, i.e. the rest of the components (displayed in black in each plot). The number of new users arriving in each compartment is shown in green. The size of the strongly connected component is approximately 10 times smaller that the size of the weakly connected component. New users arrive principally in the weakly connected giant component or the corona. The shaded areas represents important events: period: the Associated Press announcing of Clinton winning the nomination (June 6), Bernie Sanders officially terminating his campaign and endorsing Clinton (July 12) and the Republican (June 18-21) and Democratic (June 25-28) Conventions. An increase in the size of the different component occurs during these events.

in the pro-Trump (n=57) than in the pro-Hillary (n=24) categories and approximately the same number in the anti-Trump (n=36) and anti-Hillary (n=38) categories. The number of tweets using at least one of the classified hashtags amount for 32% of all the tweets having at least one hashtag.

## 2.3 Predicting election trends

The absolute number of users expressing support for Clinton and Trump as well as relative percentage of supporters to each party's candidate in the strongly connected component and in the entire population dataset is shown in Figs. 4 and 5, respectively. Results for the weakly connected component are similar to the whole population. The support of each users is assigned to the candidate for which the majority of its daily tweets are classified (see methods 5.2). Approximately 4.5% of the users are unclassified every day, as they posts the same number of tweets supporting Trump and Clinton. These users can be considered as "undecided".

We find important differences in the popularity of the candidates according to the giant components considered. The majority of users in the SCGC is generally in favor of Donald

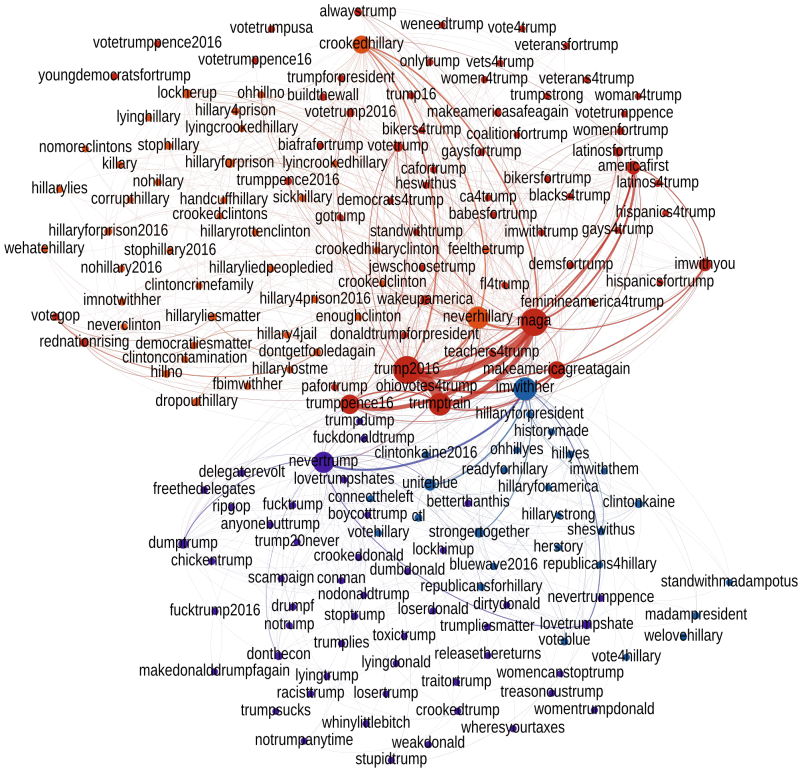


**Figure 3 | Hashtag classification.** Network of hashtag obtained by our algorithm. Nodes of the network represent hashtags and an edge is drawn between two hashtags when they appear in the same tweet. The size of the node is proportional to the total number of occurrence of the hashtag and the width of the edges is proportional to the number of times the two hashtag appeared together. The network visualization is obtained with a force-layout algorithm where nodes repulse and edges attracts. Two main clusters are visible, corresponding to the Pro-Trump/Anti-Clinton and Pro-Clinton/Anti-Trump hashtags. Inside of these two clusters, the separation between Pro-Trump (red) and Anti-Clinton (orange), or Pro-Clinton (blue) and Anti-Trump (purple), is also visible.

Trump for most of the time of observation (Fig. 4). However, the situation is reversed, with Clinton being more popular than Trump, when the entire Twitter dataset population is taken into account (Fig. 5), revealing a difference in the network localization of the users belonging to the different parties. A difference in the dynamics of the supporters opinion is also uncovered: during important events, corresponding to large spikes in the size of the network, Hillary Clinton's supporters become the majority inside the SCGC. It is as though Trump supporters dominate the campaign machinery inside the most important strong component, yet, this domination does not extend to the whole Twitter electoral population, since Clinton still has more supporters in the population at large. These Clinton supporters are not as active as the Trump supporters, except when there is a large event (like the conventions), when party bloggers may be activated.

More importantly, we next compare the daily global opinion measured in our whole dataset with the opinion obtained from traditional polls. We use the National Polling Average computed by the New York Times (NYT) [1] which is a weighted average of all polls

[1] `http://www.nytimes.com/interactive/2016/us/elections/polls.html`

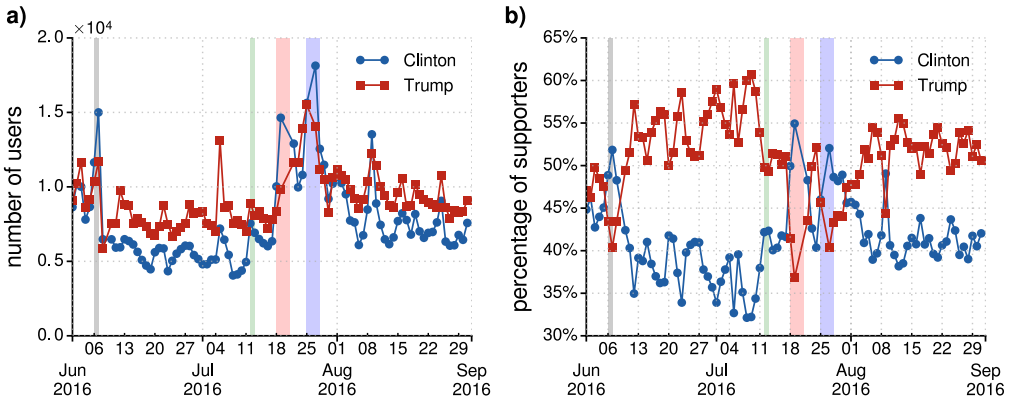


**Figure 4 | Supporters in the strongly connected giant component.** (a) Absolute number and (b) percentage of supporters of Trump (red) labeled as Pro-Trump or Anti-Clinton (red) and Clinton, labeled as Pro-Clinton or Anti-Trump (blue) inside the strongly connected giant component as a function of time. The opinion of the strongly connected giant component is in favor of Donald Trump in between important events but during important events, the opinion shifts in favor of Hillary Clinton. In (b) the data adds to 100% when considering the unclassified users ($\approx 4.5\%$).

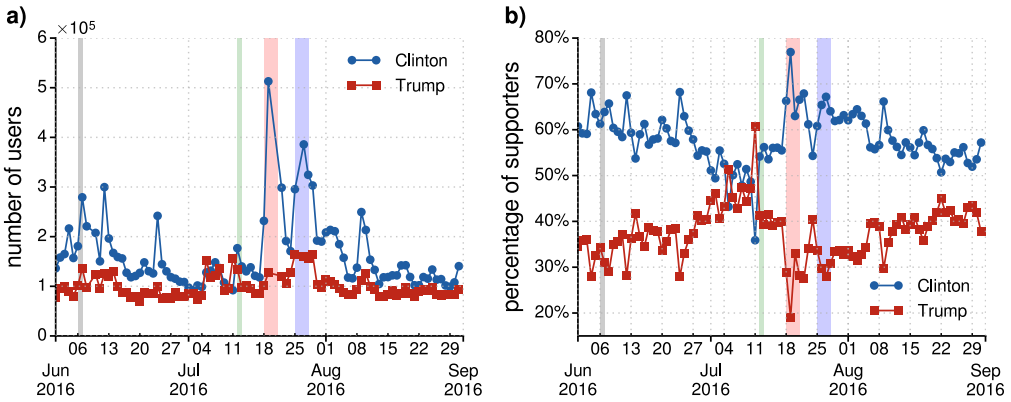


**Figure 5 | Supporters in the general Twitter election population.** (a) Total number and (b) percentage of users labeled as Pro-Trump or Anti-Clinton (red) and as Pro-Clinton or Anti-Trump (blue) in the entire Twitter population talking about the election campaign as a function of time. Taking into account all the users in our dataset, the popular opinion is generally in favor of Hillary Clinton in contrast with the strongly connected giant component in Fig. 4. The popularity of Donald Trump increases before the Conventions. During the Conventions a large drop in his popularity is visible. His popularity slowly increases again after the Conventions, but Hillary Clinton still dominates the opinion. According to this result, as of the final date of acquisition on September 1st, 2016, Clinton is leading the popular vote. Updates will be provided at `kcorelab.com`.

(total 270) listed in the Huffington Post Pollster API [2]. Greater weight are given to polls conducted more recently and polls with a larger sample size. Three types of traditional polls are used: live telephone polls, online polls and interactive voice response polls. The sample size of each polls typically varies between several hundreds to tens of thousands respondents and therefore the aggregate of all polls considered by the NYT represents a sampling size in the hundred of thousand of respondents.

[2] `http://elections.huffingtonpost.com/pollster/api`

The comparison between our Twitter prediction and the New York Times national polling average is shown in Fig. 6. The global opinion obtained from our Twitter dataset is in excellent agreement with the NYT polling average, giving the majority to Hillary Clinton. The scale of the oscillations visible in the supporter trends in Twitter and in the NYT polls are also in agreement beyond the small scale fluctuations which are visible in the Twitter opinion time series since it represents a largely fluctuating daily average. Furthermore, a time shift is apparent between the opinion in Twitter and in the NYT polls in the sense that the Twitter data anticipates the NYT National Polls by several days. This shift reflects the fact that Twitter represents the fresh instantaneous opinion of its users while traditional polls may represent a delayed response of the general population that takes more time to spread, as well as typical delays in performing and compiling traditional polls by pollsters.

In order to precisely evaluate the agreement between Twitter and NYT time series, we perform a least square fit of a linear function, followed by a moving average, of the Twitter popularity percentage of supporters of each candidate to their NYT popularity percentage. Specifically, we apply the following transformation:

$$r'_i(t) \mapsto A\ r_i(t - t_d)\ +\ b, \tag{1}$$

where $r_i(t)$ is the ratio of users in favor of candidate $i$ =\{Trump, Clinton\} at time $t$ and we perform a moving average over a time window of $w$ days. The moving window average is done to convert fluctuating daily data into a smooth trend that can be compared with the NYT smooth time series aggregated over many polls. However, the raw daily data remains as the significant prediction from Twitter data.

The constants $A$ and $b$ represent rescaling parameters that fit the actual percentage of the NYT polls. It is important to note that Twitter data cannot predict the exact percentage of supporters to each candidate in the general population due to the uncertainty about the number of voters that do not express their opinion on Twitter and about the number of users that are undecided. However, is more important to capture the relative trend of both candidates popularity respect to each other, which is obtained from Twitter. Even if Twitter may not provide the exact percentage of support for each candidate nationwide, the relevant relative opinion trend is fully captured by Twitter with high precision. Furthermore, the important parameter is $t_d$, the time delay between the anticipated opinion trend in Twitter and the delayed response captured by the NYT population at large. This delay time is independent from the actual value of the average popularity of each candidate.

The constant parameters that provide the best fit between the Twitter data and NYT in the least square sense are: $A = 0.21$, $b = 0.32$, $t_d = 6$ days for Hillary Clinton's election popularity and $A = 0.22$, $b = 0.31$, $t_d = 15$ days, for Donald Trump's election popularity (window average is $w = 13$ days for both candidates). The results of the fit, before and after the moving average, are shown in Fig 6. The Pearson product-moment correlation coefficients for each fit have a remarkable high value, $r = 0.9$, for Trump and Clinton data showing the high accuracy of the agreement between our predictions and the NYT aggregate. Our results show that the instantaneous opinion on Twitter not only follows the national average but, more interestingly, it forecasts the opinion in

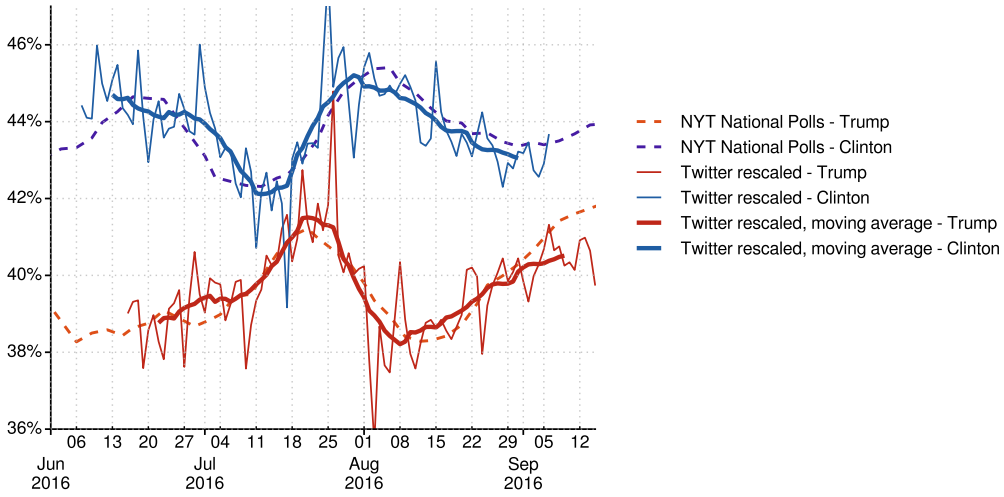


**Figure 6 | Validation of Twitter election trend and NYT aggregate national polls.** Least square fit of the percentage of Twitter supporters in favor of Donald Trump and Hillary Clinton with the results of the polls aggregated by the New York Times for the popular votes. Twitter opinion time series are in close agreement with the NYT National Polls. As Twitter provides an instantaneous measure of the opinion of its users, a time-shift exists between the New York Times polls and the Twitter opinion. Twitter opinion forecasts Donald Trump's NYT poll by 15 days and Hillary Clinton's poll by 6 days. Pearson's coefficient between the NYT and the moving averaged Twitter opinion has a remarkable high value $r$ =0.9 using an averaging window of length 13 days, for Trump and Clinton.

the NYT traditional national polls by 6 days for Hillary Clinton and 15 days for Donald Trump. Such a temporal anticipation between the instantaneous response in Twitter and the delayed response in the general polling population may become crucial when accurate instantaneous election trends are needed, specially near the general election day.

## 3 Update: last weeks of the election

We repeat our analysis to measure the opinion trend on Twitter during the last weeks of the elections: from September 26 to November 6. This period cover the three presidential debates (September 26, October 9 and October 19).
Figure 7 shows the total number and the percentage of users classified as pro-Trump and pro-Hillary supporters during the last weeks of the elections.
Figure 8 shows the percentage of users classified as pro-Trump and pro-Hillary supporters in Twitter compared with the NYT polling average fitted to a seven day moving average of the Twitter data. Twitter opinion time-series forecast the NYT average by 5 days for Donald Trump and 13 days for Hillary Clinton. The release on October 28 of letter from FBI Director James B. Comey to the Congress saying that new emails, potentially linked to the closed investigation into whether Hillary Clinton had mishandled classified information, had be found created a large temporary spike in favor of Donald Trump on the same day. However Hillary Clinton supports on Twitter is increasing again after this event.
An average of the support percentage for the last 7 days shows Clinton as the winner of the Twitter opinion by 55.5% versus Trump 44.5% (normalized to two candidates).

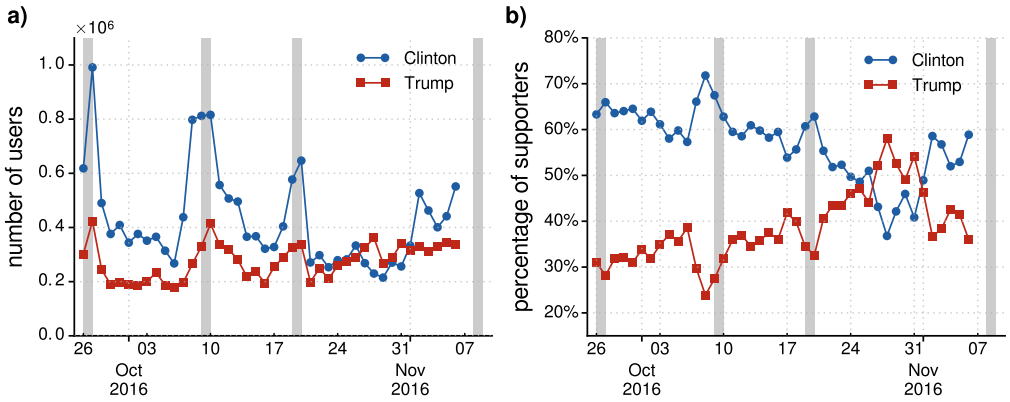


**Figure 7 | Supporters in the general Twitter election population.** (a) Total number and (b) percentage of users labeled as Pro-Trump or Anti-Clinton (red) and as Pro-Clinton or Anti-Trump (blue) in our entire Twitter dataset about the election campaign as a function of time. The three debates, represented by vertical graylines, are clearly visible with spikes in the number of total users that are mainly in favor of Hillary Clinton. A spike in the percentage of pro-Trump users is visible on October 28 and corresponds to the day that the FBI Director, James B. Comey, sent a letter to the Congress saying that new emails, potentially linked to the closed investigation into whether Hillary Clinton had mishandled classified information, had be found. Although this event correspond to an almost 60% share of total supporters on this day, looking at the total number of users (a), we see that it mainly correspond to a lack of activity from Clinton's supporters. On November 6, FBI Director James Comey told lawmakers that the FBI hasn't changed its opinion that Hillary Clinton should not face criminal charges after a review of new emails. This announcement correspond to another spike in Twitter opinion favor of Hillary Clinton. Our previous analysis is confirmed by these new results, i.e.: Hillary supporters are the majority but they mainly react when important event happen.

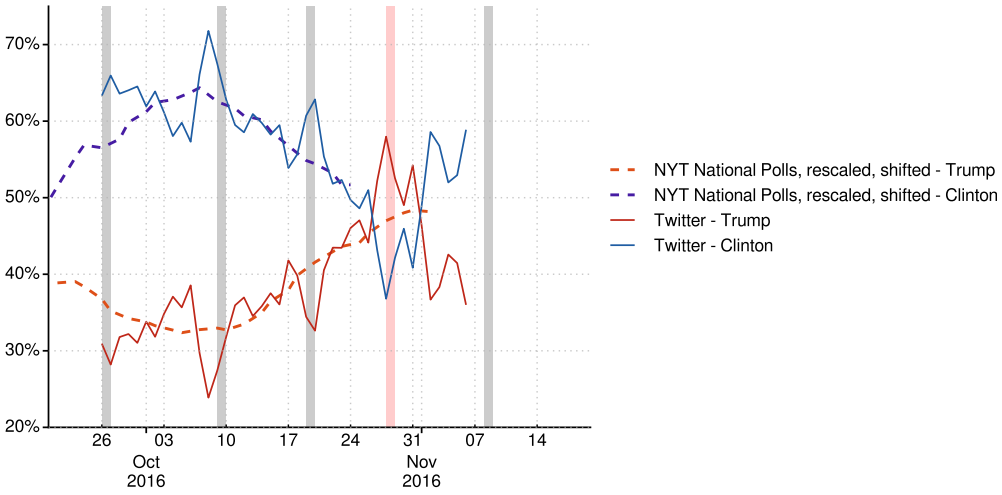


**Figure 8 |** Final prediction of the Twitter opinion before the election

# 4 Conclusions

We find a remarkable high correlation between our predicted Twitter opinion trends and the New York Times polling national average. This result validates the use of Twitter to capture global trends in the population at large. More importantly, the opinion trend in Twitter is instantaneous and anticipates the NYT aggregated surveys by 6 to 15 days,

showing that Twitter can be used as an early warning signal of global opinion trends happening in the general population at the country level.

Our results also reveal a difference in the behavior of Twitter users supporting Donald Trump and users supporting Hillary Clinton. Donald Trump's supporters are almost constantly the majority in the strongly connected giant component, revealing their higher involvement in reciprocal interactions. We also discover a larger number of hashtags expressing support for Donald Trump than Hillary Clinton, suggesting the fact that Trump supporters are more active and prone to create and share hashtags (including a larger number of tweets with hashtags pro trump). Trump supporters are also more compromised and opinionated than Clinton supporters as evidenced by their larger activity in the strongly connected component. While these evidences suggest that Donald Trump's supporters may be dominating the Twitter-sphere, the analysis of the dynamics of the whole network reveals a more contrasted picture. Indeed, during important events, such as when the Associated Press declared that Hillary Clinton would be the Democratic Nominee and during the two National Conventions, the opinion inside the strongly giant component shifts in favor of Hillary Clinton. These events also correspond to large positive fluctuations in the network size. This suggests that Clinton's supporters express their opinion more strongly during important events and are less active than Trump's supporters the rest of the time.

While Trump supporters might be gaining the race inside the strongly connected giant component, thus forming a very cohesive group with large influence at the core of the network, Clinton still wins the popular vote when we consider all the Twitter supporters, and not only those in the SCGC. Thus, while the Twitter campaign of Clinton seems to be less enthusiastic and dormant compared to the Twitter Trump machinery, candidate Clinton still wins the popular vote in the whole network. Following the adagio "one man, one vote", the important result remains the Clinton's prevalence in the whole network, not just the SCGC. Yet, it remains a question whether the results at the whole Twitter population level will be reflected during the real election day with the victory of Clinton as predicted, or, in contrast, the opinion of the SCGC will prevail giving rise to a victorious Trump campaign.

In conclusion, we present a Twitter analytics platform (available at `kcorelab.com` where updates will be posted) that uses a combination of percolation and statistical physics of networks, natural language processing and machine learning to uncover the opinion of Twitter users. The remarkable agreement between our predictions and the national polls indicate that Twitter analytics can be used to predict global election trends of the large-scale population. Our results validate the use of our Twitter analytics machinery as a mean of assessing opinions about political elections, which, indeed, comes at a fraction of the cost of traditional NYT polling methods employed by aggregating the whole of the US$ 18 billion-revenue market research and public opinion polling industry (NAICS 54191).

In contrast to traditional unaggregated polling campaigns which are unscalable, typically ranging at most in the few thousand respondents, our techniques have the advantage of being highly scalable as they are only limited by the size of the underlying exponentially-growing social networks. Moreover, traditional polling suffers from a declining rate of respondents being only 9% according to current estimates (2012) down from 36% in 1997

[3], while current social media grow in the billion of people population wide. Although the demographics representation of Twitter is not perfect [4] traditional polls also suffer from large margin of error [5]. In a recent NYT-The Upshot study, it was found that four pollsters arrived at different estimates even if when they use the same raw data from the respondents [6]. This uncertainty in the individual polls comes from the judgment of pollster to adjust their raw data samples to match the demographics of the electorate. As different pollsters rely on different weighting methods, the results among polls vary well above the reported "margin of error" (error due to sampling which represents the likehood that the result from the sample is close to the whole population result, and it is usually reported in a range around $\pm$ 3%). Thus, in general, average polls over a large number of pollsters (like the NYT aggregate) may be needed for accurate forecasting rather than individual polls. While large-scale aggregates of hundred of polls can be obtained for important events like the general elections, tracking social trends in general cannot rely on such expensive aggregates which use the combined power of the whole multibillion-dollar polling industry. In such cases, mining the opinion from social media outlets might be the only option for accurate forecasting.

As online social media usage continue to grow, our trend-predictor machinery may become closer and closer to the true opinion of the whole population, and may, perhaps, end up rendering traditional polling methods obsolete in the not too distant future. A further advantage is that, as opposed to traditional opinion polls where people are asked directly a series of questions, our analytics infer the opinion of people from their writings and tweetings without direct interaction with the respondents. The methods can be extended to assess any kind of trend from social media, ranging from the opinion of users regarding products and brands, to political movements, thus, unlocking the power of Twitter to understand trends in the society at large.

# 5 Methods

## 5.1 Data collection and social network reconstruction

We collected tweets using the Twitter Search API from June 1st, 2016 to September 1st, 2016. We gather a total of 73 million tweets mentioning the two top candidates from the Republican Party (Donald J. Trump) and Democratic Party (Hillary Clinton) by using two different queries with the following keywords: *trump OR realdonaldtrump OR donaldtrump* and *hillary OR clinton OR hillaryclinton*. For every day in our dataset, we construct the social network $G(V, E)$ where $V$ is the set of vertices representing users and $E$ is the set of edges representing interactions between the users. In this network, edges are directed and represent influence. When a user $v_i \in V$, retweets, replies to, mentions or quotes an other users $v_j \in V$, a directed edge is drawn from $v_j$ to $v_i$. We remove Donald Trump

[3] http://www.people-press.org/2012/05/15/assessing-the-representativeness-of-public-opinion-surveys

[4] http://www.pewinternet.org/2015/08/19/the-demographics-of-social-media-users/

[5] http://www.nytimes.com/2016/10/06/upshot/when-you-hear-the-margin-of-error-is-plus-or-minus-3-percent-th
html

[6] http://www.nytimes.com/interactive/2016/09/20/upshot/the-error-the-polling-world-rarely-talks-about.
html

| F1 | AUROC | Accuracy | Precision | Recall |
|---|---|---|---|---|
| 0.81 | 0.89 | 0.81 | 0.82 | 0.80 |

**Table 1 |** Best classification score over 10-fold cross-validation, achieved using a Logistic Regression Classifier with $L_2$ regularization.

(*@realdonaldtrump*) and Hillary Clinton (*@hillaryclinton*) from the network, as we are interested by the opinion and dynamics of the rest of the network. We divide the network in three compartments: the strongly connected giant component (SCGC), the weakly connected giant component (WCGC) and the corona (Fig. 2). The SCGC is defined as the largest maximal set of nodes where there exists a path in both directions between each pair of nodes. The SCGC is formed by the central, most densely connected region of the network where the influencers are located, and where the interactions between users are numerous. The WCGC is the largest maximal set of nodes where there exists a path in at least one direction between each pair of nodes. The corona is formed by the smaller components of remaining users and the users that were only connected to Hillary Clinton or Donald Trump official accounts, which were removed for consistency.

### 5.2 Opinion mining

We build a training set of labeled tweets with two classes: 1) pro-Clinton or anti-Trump and, 2) pro-Trump or anti-Clinton. We discard tweets belonging to the two classes simultaneously to avoid ambiguous tweets. We also remove retweets to avoid duplicates in our training set. We also select only tweets that were posted using an official Twitter client in order to discard tweets that might originate from bots and to limit the number of tweets posted from professional accounts. We use a balanced set, with the same number of tweets in each class, totaling 798,354 tweets. The tweet contents is tokenized to extract a list of words, hashtags, usernames, emoticons and urls. We test the performance of different classifiers (Support Vector Machine, Logistic Regression and modified Huber) with different regularization methods (Ridge Regression, Lasso and Elastic net). Hyperparameter optimization is performed with a 10-fold cross validation optimizing $F_1$ score. The best score is obtained with a Logistic Regression classifier with $L_2$ penalty (Ridge Regression). Classification scores are summarized in Table 1.

## References


[1] A. Mislove, S. Lehmann, Y.-Y. Ahn, J.-P. Onnela, and J. N. Rosenquist, Pulse of the nation: Us mood throughout the day inferred from twitter, 2010.

[2] A. Hannak *et al.*, Proc. of the 6th International AAAI Conference on Weblogs and Social Media , 479 (2012).

[3] A. Pak and P. Paroubek, Twitter as a Corpus for Sentiment Analysis and Opinion Mining Microblogging Microblogging = posting small blog entries Platforms, in *Proceedings of the Seventh conference on International Language Resources and Evaluation*, pp. pp. 19–21, Valletta, Malta, 2010.

[4] W. Quattrociocchi, G. Caldarelli, and A. Scala, Scientific Reports **4**, 4938 (2014).

[5] N. F. Johnson *et al.*, Science **352**, 1459 (2016).

[6] Y. Liu, X. Huang, A. An, and X. Yu, ARSA, in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval - SIGIR '07*, p. 607, New York, New York, USA, 2007, ACM Press.

[7] J. Bollen, H. Mao, and X. Zeng, Journal of Computational Science **2**, 1 (2011), 1010.3003.

[8] G. Ranco, D. Aleksovski, G. Caldarelli, M. Grčar, and I. Mozetič, PLOS ONE **10**, e0138441 (2015), 1506.02431v2.

[9] B. O'Connor, R. Balasubramanyan, B. R. Routledge, and N. A. Smith, Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media From , 122 (2010).

[10] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe, Social Science Computer Review **29**, 402 (2011).

[11] J. Borondo, A. J. Morales, J. C. Losada, and R. M. Benito, Chaos: An Interdisciplinary Journal of Nonlinear Science **22**, 023138 (2012), 1309.5014.

[12] D. Gayo-Avello, Social Science Computer Review **31**, 649 (2013), 1206.5851.

[13] G. Caldarelli *et al.*, PLoS ONE **9**, e95809 (2014).

[14] V. Kagan, A. Stevens, and V. S. Subrahmanian, IEEE Intelligent Systems **30**, 2 (2015).

[15] J. Serrano-Guerrero, J. A. Olivas, F. P. Romero, and E. Herrera-Viedma, Information Sciences **311**, 18 (2015), arXiv:1201.4597v1.

[16] M. M. Bradley and P. P. J. Lang, Psychology **Technical**, 44 (1999), arXiv:1011.1669v3.

[17] V. Subrahmanian and D. Reforgiato, IEEE Intelligent Systems **23**, 43 (2008).

[18] A. Montejo-Ráez, E. Martínez-Cámara, M. T. Martín-Valdivia, and L. A. Ureña-López, Computer Speech & Language **28**, 93 (2014).

[19] Y. R. Tausczik and J. W. Pennebaker, Journal of Language and Social Psychology **29**, 24 (2010).

[20] A. Go, R. Bhayani, and L. Huang, Technical report **150**, 1 (2009).

[21] S. Pei, L. Muchnik, J. S. Andrade, Jr., Z. Zheng, and H. A. Makse, Scientific Reports **4**, 5547 (2014), 1405.1790.

[22] A. Bunde and S. Havlin, *Fractals and Disordered Systems* (Springer Berlin Heidelberg, 2012).

[23] B. Bollobás, *Random Graphs*Cambridge Studies in Advanced Mathematics (Cambridge University Press, 2001).